



AFRL-AFOSR-JP-TR-2017-0075

Autonomous Learning in Mobile Cognitive Machines

Byoung-Tak Zhang
SEOUL NATIONAL UNIVERSITY

11/25/2017
Final Report

DISTRIBUTION A: Distribution approved for public release.

Air Force Research Laboratory
AF Office Of Scientific Research (AFOSR)/ IOA
Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Executive Services, Directorate (0704-0188). Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.</p>					
1. REPORT DATE (DD-MM-YYYY) 27-11-2017		2. REPORT TYPE Final		3. DATES COVERED (From - To) 26 Aug 2016 to 25 Aug 2017	
4. TITLE AND SUBTITLE Autonomous Learning in Mobile Cognitive Machines			5a. CONTRACT NUMBER 5b. GRANT NUMBER FA2386-16-1-4089 5c. PROGRAM ELEMENT NUMBER 61102F		
6. AUTHOR(S) Byoung-Tak Zhang			5d. PROJECT NUMBER 5e. TASK NUMBER 5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) SEOUL NATIONAL UNIVERSITY SNUR&DB FOUNDATION RESEARCH PARK CENTER SEOUL, 151742 KR				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AOARD UNIT 45002 APO AP 96338-5002				10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/AFOSR IOA 11. SPONSOR/MONITOR'S REPORT NUMBER(S) AFRL-AFOSR-JP-TR-2017-0075	
12. DISTRIBUTION/AVAILABILITY STATEMENT A DISTRIBUTION UNLIMITED: PB Public Release					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT <p>Intelligence is a capability ascribed typically to animals, but not usually to plants. Animals can move while plants do not. Is the mobility a necessary condition or driving force for the emergence of intelligence? The researchers hypothesize that mobility plays a foundational role in evolving animal and human intelligence, thus, is fundamentally important in understanding and creating embodied cognitive systems. In this project, the researchers aim to develop a new class of machine learning algorithms for mobile cognitive systems that actively collect data by sensing and interacting with the environment. They envision a new paradigm of autonomous AI that overcomes the previous AI paradigms of top-down/rule-driven symbolic and bottom-up/data-driven statistical systems. Inspired by the dual process theory of mind. They use mobile robot platforms to investigate the autonomous learning algorithms and demonstrate their capability in real-world home environments. The hypothesis of the brain being evolved to support its mobility has been raised. In fact, as the project progressed, the researchers discovered that if one of the perception-action-learning is missing or malfunctioning, maintaining the full ability of the robot was almost impossible in functioning in given scenarios. However, the researchers believe that even though perception is very important, if it is unable to perform actions in the environment, the perception ability almost loses its purpose for mobile robots in a home environment. In the basic year of this project, the researchers achieved a basic system for mobile robots to perceive, act and learn within the environment. They believe that using this system as a base, developing higher functions like memory and planning could be attained, which would be a significant step forward to achieving a truly human-level AI.</p>					
15. SUBJECT TERMS Autonomous Agents, Learning Algorithms, Cognitive machines					
16. SECURITY CLASSIFICATION OF: a. REPORT Unclassified			17. LIMITATION OF ABSTRACT SAR	18. NUMBER OF PAGES 11	19a. NAME OF RESPONSIBLE PERSON ROBERTSON, SCOTT 19b. TELEPHONE NUMBER (Include area code) +81-042-511-7008

“Autonomous Learning in Mobile Cognitive Machines”

2017. 11. 25

Name of Principal Investigators (PI and Co-PIs):

- E-mail Address: btzhang@bi.snu.ac.kr
- Institution: Seoul National University
- Mailing Address: ROOM 417, BLD. 138, SEOUL NATIONAL UNIVERSITY 599 GWANAK-RO, GWANAK-GU, SEOUL 151742 KOREA, REPUBLIC OF
- Phone : +82-10-8647-7381
- Fax : +82-2-875-2240

Period of Performance: 08/26/16– 08/25/17

Abstract: Intelligence is a capability ascribed typically to animals, but not usually to plants. Animals can move while plants do not. Is the mobility a necessary condition or driving force for the emergence of intelligence? We hypothesize that mobility plays a foundational role in evolving animal and human intelligence, thus, is fundamentally important in understanding and creating embodied cognitive systems [1]. In this project, we aim to develop a new class of machine learning algorithms for mobile cognitive systems that actively collect data by sensing and interacting with the environment. We envision a new paradigm of autonomous AI that overcomes the previous AI paradigms of top-down/rule-driven symbolic and bottom-up/data-driven statistical systems. Inspired by the dual process theory of mind [2]. We use mobile robot platforms to investigate the autonomous learning algorithms and demonstrate their capability in real-world home environments.

Introduction: In the history of artificial intelligence (AI), two main approaches have emerged: symbolic and statistical systems. The former approach, or first generation AI, is deductive, relies on rule-based programming, and can solve complex problems, however, faces difficulties in learning and adaptability. The latter approach, or second generation AI, is inductive, relies on statistical learning from big data, but cannot solve complex problems, the speed of learning is limited, and thus faces the issues of scalability. To create human-level artificial intelligence, we need a methodology that combines the best of both approaches and also scales up to real complex problems.

Recent advancements in deep learning provide a crucial lesson in this direction, i.e., building more expressive representations help solve complex problems [3][4]. This provides evidence for an earlier prediction, that “learning requires much more memory than we have thought to solve real-world problems” [5]. Deep learning models use much larger memory than previous machine learning models, but they do not overfit due to the increased data size. However, deep learning models are very limited in their learning speed, flexibility, and robustness when applied to dynamic environments of mobile cognitive agents.

Why and how has the human brain evolved to learn so rapidly, flexibly, and robustly? We hypothesize that the brain evolved these properties mainly to support its mobility for the survival of its body in hostile environments [1][6]. In fact, the brain’s main function is to make decisions and control the body motion. Higher functions like memory and planning were evolved on top of this substrate. Therefore, to achieve a truly human-level AI, it is important to study higher-level intelligence, such as vision and language, in a mobile platform and dynamic environment. It is our belief that fast, flexible, and robust learning in interactive mobile environments will give rise to a new paradigm of machine learning that will enable the next generation of autonomous AI systems.

In this project, the ultimate goal is to demonstrate a mobile personal robot that learns the objects, people, actions, events, episodes and schedule plans from daily to extended periods of time. In the basic year of the project, we built a multi-module integrated system for mobile robots to perceive information (objects, people, actions) from the environment, act (schedule, interact) according to the perceived information and develop models that learn the dynamics of the environment. We also demonstrated the integration of multimodal information for an interactive system which efficiently infers and responds to the goals and plans of the observed environment.

Experiments and Results:

a) Perception-Action-Learning System for Mobile Social-Service Robots

Making robots becoming more human-like, capable of providing natural social services to the customers in dynamic environments such as houses, restaurants, hotels and even airports has been a challenging goal for researchers in the field of social-service robotics. One promising approach is developing an integrated system of methodologies from many different research areas. This multi-module integrated intelligent robotic system has been widely accepted and its performance has been well known from previous studies [7][8]. However, with the individual roles of each module in the integrated system, perception modules mostly suffered from desynchronization between each other and difficulty in adapting to dynamic environments [9]. This occurred because of the different process time and scale of coverage of the adopted vision techniques [10]. To overcome such difficulties, developers usually upgraded or added expensive sensors (hardware) to the robot to improve performances. Though this may have provided some solutions to the limitations, current robot systems still have difficulties on natural interaction within real-life, dynamic environment.

We account this matter by designing a system incorporated with state-of-the-art deep learning methods and inspiration by the cognitive perception-action-learning cycle [11]. The implemented novel and robust integrated system for mobile social-service robots that at least includes an RGB-D camera and any obstacle detecting sensors (laser, bumper, sonar), achieved real-time performance on various social service tasks. Also, by performing the task in real-time with robustness, more natural interaction with people could be attained.




Figure 1. Perception-Action-Learning system for mobile social-service robots using deep learning

As illustrated in Figure 1, our system's perception-action-learning cycle works in real-time (~0.2 s/cycle) where the arrows indicate the flow of each module. The system was implemented on a server of I7 CPU, 32 GB RAM and GTX Titan 12 GB GPU. Using ROS topics, the communication between the server and the robot were achieved and the ROS topics were passed through 5 GHz Wi-Fi connection.

The conducted experiments were finely designed by the RoboCup@Home Committee, which is described in the rulebook [12] and our system was able to perform all the scenarios in a significantly improved way.

RoboCup2017@Home Social Standard Platform League (SSPL) Winning First Place

We used our system on SoftBank Pepper, a standardized mobile social-service robot, and achieved the highest score in every scenario performed at the RoboCup2017@Home Social

Standard Platform League (SSPL), winning first place overall.

Our system allows robots to perform social service tasks in real-life social situations with high performance working in real-time. However, our system is yet to fulfill every individual's expectations on performance and processing speed, we highlight the importance of research on not only the individual elements but the integration of each module for developing a more human-like, idealistic robot to assist humans in the future. Related videos can be found at <https://goo.gl/Pxnf1n> and our open-sourced codes at https://github.com/soseazi/pal_pepper.

[Table 1] RoboCup2017@Home Social Standard Platform League (SSPL) Test 1 Result

Team	Poster	Speech Person &	Cocktail Party	Help Me Carry	GPSR	Total	Rank
AUPAIR	45.00	117.5	30	10	42.5	245.00	1
UTS Unleashed	33.33	85.5	27.5	0	17.5	163.83	2
SPQReL	41.67	32.5	10	5	7.5	96.67	3
KameRider	31.67	60	0	0	0	91.67	4
UChilePeppers	31.67	50	0	0	0	81.67	5
UvA@Home	20.00	47.5	0	0	0	67.50	6
ToBI@Pepper	41.25	17.5	7.5	0	0	66.25	7

[Table 2] RoboCup2017@Home Social Standard Platform League (SSPL) Test 2 Result

Team	Stage 1	Open Challenge	Tour Guide	Restaurant	EE-GPSR	Total	Rank
AUPAIR	245.00	178.47	95	40	70	628.47	1
UTS Unleashed	163.83	121.53	0	0	0	285.36	2
SPQReL	96.67	130.56	0	10	20	257.22	3
KameRider	91.67	136.81	0	15	0	243.47	4

b) Integrated Perception Towards Fully Autonomous General Purpose Service Robots

To interact with or assist people, service robots require a perception framework that can provide information such as the location/type of objects and the identity/pose/gender of people in the environment. Many perception frameworks have been used in service robots. OpenCV or OpenNI have been widely used to perform perception tasks such as object detection, human pose estimation. These frameworks focus on only a few tasks such as object detection or face recognition. Furthermore, these frameworks use traditional vision methods that are known to be vulnerable to illumination change or translation of objects. Those frameworks also lack a reasoning engine that can build perceptual information and reason about it. Frameworks such as RoboSherlock [13][14] provide sophisticated reasoning engines on top of the integrated perception pipeline but they focus only on object manipulation and they also use traditional vision modules. These limitations in perception frameworks often limit service robots to only show good performance in well-defined tasks in a controlled environment.

Recently, following the remarkable success of deep learning in object recognition [15], many deep learning based perception models have been proposed. Deep learning based approaches are known to be robust to illumination or translation and have marked state-of-the-art performances in many vision tasks such as object detection [16][17][18], image description [19][20], and pose estimation [21][22]. These models show superior performance than more traditional approaches. However, these models are not enough to be deployed in complex and realistic perception tasks since they mostly focus on individual tasks such as object detection, face detection, or object recognition. Furthermore, these models also lack reasoning engines that can process perceptual information efficiently.

We propose IPSRO (Integrated Perception for Service RObots) framework, which is ROS-friendly integrated perception system that we have recently open-sourced. IPSRO can flexibly integrate several perception modules including deep learning models to extract rich and

useful perceptual information from the environment based on a unified perception representation. On top of that, IPSRO can process the generated perceptual information to perform complex perception tasks. We conducted experiments using GPSR (General Purpose Service Robot) task of RoboCup@Home. In the GPSR task, the robot has to execute arbitrary voice commands. The commands include but are not limited to finding an object or person, answering a question, following or guiding a person, counting the number of objects, and describing a person or place.

GPSR (General Purpose Service Robot) Command Executing Experiments




Figure 2. Visualization of perceptual information of IPSRO framework.

(Command 1, 2, 3 from left to right)

We conducted experiments using the GPSR task in a lab environment and RoboCup@Home social standard platform league. In the lab environment, we gave three following commands to the robot with IPSRO framework.

1. Find James in the kitchen and guide him to entrance
2. Tell me how many coke bottles are on the table
3. Describe unknown people in the living room
in Figure 2 of page 1 of the following

As seen in Figure 2, our framework successfully extracted all relevant tags from the camera image and succeeded to execute all commands correctly.

In the RoboCup@Home competition, the robot with our framework is given three following commands

1. Say the time to Jacob at the kitchen table
2. Find a hairspray at the kitchen table
3. Tell me how many cokes are on the desk

Our robot successfully executed all commands, scoring highest score among seven teams (Table 1). The video of the GPSR competition can be found in goo.gl/fyRhtD

c) Robust Human Following by Deep Bayesian Trajectory Prediction for Home Service Robots




Figure 3. Robust following for robots to interact with people in a home environment

Human following by a robot has been an ongoing research topic in the robotics community [23], with annual robotic competitions [12] [24] to test the following performances. To achieve such an ability, previous studies worked with vision techniques to capture the human's characteristic features to detect and track the human. For example, SIFT [25], ORB [26] and template matching [27] were used in human tracking. However, these approaches had several limitations with in illumination change, translation of objects and occlusion of the sensors. Moreover, the difficulty of separating a person between the foreground and the background was a very demanding issue to maintain a following system with a certain level of performance.


In contrast with the mentioned literature, combining the high performance in recognition using deep learning methods, empowered by the computational power of GPU, and generally adoptable ROS system, we introduce a robust integrated system for home service robots to follow a person in the home environment called Deep Bayesian Trajectory Prediction (DBTP). DBTP contributes with 1) robust detection and identification of a person in real-time (around 0.3 s) in a homelike environment with state-of-the-art performance, 2) following the target with contextual information to perform better collision avoidance and 3) by recording a person's coordinate trajectory in real-time matter, we could empower the robot with an ability to follow the person with variational Bayesian linear regression (VBLR) [28] based trajectory prediction when the robot failed to continuously follow or lost the target person it was following.

We have designed four experiments to demonstrate the proposed framework's success in following the target person, avoiding collision and continuously following when the target person is lost, in a difficult situation in the environment (Figure 4). Lastly, we report the results of our performance with this framework in the RoboCup@Home2017 following tasks.



Figure 4. Difficult situation for robot to follow. A, B: lost target; C: wall in between

1. Following performance result:






Figure 5. a) Following the whole trajectory of the target person. b) Distance between robot and target person. The number indicates the step of following the target person

Figure 5 indicates the robot's trajectory. The blue dot (robot position) is consistently following the person even when the person changes speed and direction. Moreover, at the dotted square X, Y, Z, the target person behaves with dynamic movements like wiggling side by side, moving in a narrow space and even moving toward the robot and going pass the robot. However, our system robustly follows the target person within 2.5m distance.

2. Collision avoidance result:




Figure 6. Collision avoidance. Red box: close trajectory of target. Blue box: target going over the obstacle. Robot robustly following with reflex control

To test whether our system could perform collision avoidance when following, we placed obstacles in the environment as depicted in Figure 6. First, for the red box, the target person passes the obstacle very closely and quickly where the obstacles overlapped the person's trajectory. In this case, the control system executed the dynamics control with the reflex module together to avoid the obstacle. The blue box obstacle in Figure 6 was tested to see whether our action controller could avoid difficult situations of colliding with the obstacle. When the person went over the obstacle, it resulted in the obstacle being placed between the robot and the target person. For such a case, it is impossible for the robot to follow the target with only the dynamics control. However, our navigation control planned the path periodically in respect to the person's distance and applied the reflex module when it approached close to the obstacle, resulting in the completion of following the person to the end.

3. Recovering following when lost target:




Figure 7. Left: Difficult situation for the robot to follow. Right: Predicted trajectory result of the target using momentum, maximum likelihood, and proposed variational Bayesian linear regression (VBLR). The blue line indicates the trajectory history of the target. Red X indicates the current coordinate of the target.

We examined our methods with two difficult situations where the robot could easily lose the target person (Figure 7). Task A is a situation when the person goes out the door and immediately turns right. This made our perception module capture the target person with a slight view in between the doorway (Figure 7 solid lined box top row). Task B is when robot totally loses the perception of the target person, when the target person hides behind the wall by turning left (Figure 7 dotted lined box bottom row). We compared our proposed VBLR with two other methods.

First, with task A, every method found the target. However, the gaps between each method were large in which our method achieved almost real-time re-following at that given situation. Moreover, for task B, the other two methods failed on detection of the target person. For the momentum method, the robot was unable to move out of the doorway. The ML method predicted the trajectory to go outside but went too far to recognize the target person. As a result, even for this task, our VBLR succeeded in going out of the doorway and finding the target within an average of 3 seconds. The average consumption time with 100 trials.

The video of the DBTP can be found in <https://youtu.be/F6211GhrbbE>

Conclusion:

The hypothesis of the brain being evolved to support its mobility has been raised. In fact, as the project progressed, we could discover that if one of the perception-action-learning is missing or malfunctioning, maintaining the full ability of the robot was almost impossible in functioning a given scenarios. However, we believe that even though perception is very important, if it is unable to perform actions in the environment, the perception ability almost loses its purpose for mobile robots in a home environment. Therefore, as in the basic year of this project, we achieved the basic system for mobile robots to perceive, act and learn within the environment. We believe that using this system as a base, higher developing higher functions like memory and planning could be attained, which by stepping a bit forward to achieving a truly human-level AI.

List of Publications and Significant Collaborations that resulted from this AOARD supported project:

- a) Papers published in peer-reviewed conference proceedings:
 - Beom-Jin Lee, Jinyoung Choi, Chung-Yeon Lee, Kyung-Wha Park, Sungjun Choi, Cheolho Han, Dong-Sig Han, Christina Baek, Patrick Emaase, Byoung-Tak Zhang. "Perception-Action-Learning System for Mobile Social-Service Robots using Deep Learning." AAAI 2018 Demonstration Track. (Accepted, to be published)
- b) Manuscripts submitted but not yet published:
 - Beom-Jin Lee, Jinyoung Choi, Christina Baek and Byoung-Tak Zhang. "Robust Human Following by Deep Bayesian Trajectory Prediction for Home Service Robots." ICRA 2017 (Submitted)
 - Jinyoung Choi, Beom-Jin Lee, Chung-Yeon Lee, Kyung-Wha Park, Dong-Sig Han, Christina Baek and Byoung-Tak Zhang. "Integrated Perception Towards Fully Autonomous General Purpose Service Robots." ICRA 2017 (Submitted)
- c) Workshop

- J. Choi, B.-J. Lee, and B.-T. Zhang. "Multi-focus attention network for efficient deep reinforcement learning." AAAI 2017 Workshop on What's next for AI in games (WNAIG 2017), 2017
- d) Domestic Papers
 - S. Son, J. Kim, B.-T. Zhang. "Active Image Learning of Household Robots Using Bayesian Neural Network." Korean Institute of Information Scientists and Engineer, Winter Conference, pp. 690-692, 2016.12.
 - J. Kim. "Talking to Teach a Personal Service Robot to Get Acquainted with the Dynamically Changing Home Environment." In 2017 Korea society for Cognitive science, Annual Conference. 2017.05. (poster)
 - B.-J. Lee, J. Choi, B.-T. Zhang. "Teaching Robot to Follow a Person using Deep Reinforcement Learning." In 2017 Korea society for Cognitive science, Annual Conference, 2017.05. (poster)
 - S. Son. "Optimizing the Continual Learning of Bayesian Neural Network." In 2017 Korea society for Cognitive science, Annual Conference, 2017.05. (poster)
- e) Award
 - RoboCup@Home 2017 Social Standard Platform League 1st place
 - Video link: goo.gl/fyRhtD
- f) Open Source
 - <https://github.com/gliese581gg/IPSRO>
 - https://github.com/soseazi/pal_pepper
- g) Press
 - [2017.08.01] Seoul National University, Professor Zhang's Team, Win the 2017 National RoboCup League (<http://news.mk.co.kr/newsRead.php?year=2017&no=514908>)
 - [2017.08.07] 'AUPAIR' Winning the National Competition. Currently it is a baby step, but the potential is infinite (<http://news.joins.com/article/21823070>)

Attachments: Publications a), b) and c) listed above.

Reference:

- [1] Clark, Andy. *Being there: Putting brain, body, and world together again*. MIT press, 1998.
- [2] Kahneman, Daniel. *Thinking, fast and slow*. Macmillan, 2011.
- [3] Abadi, Martín, et al. "Tensorflow: Large-scale machine learning on heterogeneous distributed systems." *arXiv preprint arXiv:1603.04467* (2016).
- [4] Mnih, Volodymyr, et al. "Human-level control through deep reinforcement learning." *Nature* 518.7540 (2015): 529-533.
- [5] Zhang, Byoung-Tak. "Teaching an Agent by Playing a Multimodal Memory Game: Challenges for Machine Learners and Human Teachers." *AAAI Spring Symposium: Agents that Learn from Human Teachers*. 2009.
- [6] Ballard, Dana H. "Animate vision." *Artificial intelligence* 48.1 (1991): 57-86.
- [7] Brooks, Rodney. "A robust layered control system for a mobile robot." *IEEE journal on robotics and automation* 2.1 (1986): 14-23.
- [8] Siepmann, Frederic, et al. "Deploying a modeling framework for reusable robot behavior to enable informed strategies for domestic service robots." *Robotics and Autonomous Systems* 62.5 (2014): 619-631.
- [9] Rodríguez, Francisco J., et al. "A Motivational Architecture to Create more Human-Acceptable Assistive Robots for Robotics Competitions." *Autonomous Robot Systems and Competitions (ICARSC)*, 2016 International Conference on. IEEE, 2016.
- [10] Collet, Alvaro, Manuel Martinez, and Siddhartha S. Srinivasa. "The MOPED framework: Object recognition and pose estimation for manipulation." *The International Journal of Robotics Research* 30.10 (2011): 1284-1306.
- [11] Badre, David. "Cognitive control, hierarchy, and the rostro-caudal organization of the frontal lobes." *Trends in cognitive sciences* 12.5 (2008): 193-200.
- [12] B. Loy van, H. Dirk, M. Mauricio, R. Caleb, and W. Sven. (2017). Robocup@home 2017: Rules and regulations, [Online]. Available: http://www.robocupathome.org/rules/2017_rulebook.pdf

- [13] Beetz, Michael, et al. "Robosherlock: Unstructured information processing for robot perception." *Robotics and Automation (ICRA), 2015 IEEE International Conference on*. IEEE, 2015.
- [14] Bálint-Benczédi, Ferenc, Patrick Mania, and Michael Beetz. "Scaling perception towards autonomous object manipulation—in knowledge lies the power." *Robotics and Automation (ICRA), 2016 IEEE International Conference on*. IEEE, 2016.
- [15] Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. "Imagenet classification with deep convolutional neural networks." *Advances in neural information processing systems*. 2012.
- [16] Ren, Shaoqing, et al. "Faster R-CNN: Towards real-time object detection with region proposal networks." *Advances in neural information processing systems*. 2015.
- [17] Redmon, Joseph, and Ali Farhadi. "YOLO9000: better, faster, stronger." *arXiv preprint arXiv:1612.08242* (2016).
- [18] Liu, Wei, et al. "Ssd: Single shot multibox detector." *European conference on computer vision*. Springer, Cham, 2016.
- [19] Vinyals, Oriol, et al. "Show and tell: A neural image caption generator." *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2015.
- [20] Johnson, Justin, Andrej Karpathy, and Li Fei-Fei. "Densecap: Fully convolutional localization networks for dense captioning." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2016.
- [21] Toshev, Alexander, and Christian Szegedy. "Deeppose: Human pose estimation via deep neural networks." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2014.
- [22] Cao, Zhe, et al. "Realtime multi-person 2d pose estimation using part affinity fields." *arXiv preprint arXiv:1611.08050* (2016).
- [23] Sidenbladh, Hedvig, Danica Kragic, and Henrik I. Christensen. "A person following behaviour for a mobile robot." *Robotics and Automation, 1999. Proceedings. 1999 IEEE International Conference on*. Vol. 1. IEEE, 1999.
- [24] Wisspeintner, Thomas, et al. "RoboCup@ Home: Scientific competition and benchmarking for domestic service robots." *Interaction Studies* 10.3 (2009): 392-426.
- [25] Satake, Junji, Masaya Chiba, and Jun Miura. "A SIFT-based person identification using a distance-dependent appearance model for a person following robot." *Robotics and Biomimetics (ROBIO), 2012 IEEE International Conference on*. IEEE, 2012.
- [26] Munaro, Matteo, et al. "A feature-based approach to people re-identification using skeleton keypoints." *Robotics and Automation (ICRA), 2014 IEEE International Conference on*. IEEE, 2014.
- [27] Ess, Andreas, et al. "A mobile vision system for robust multi-person tracking." *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*. IEEE, 2008.
- [28] Drugowitsch, Jan. "Variational Bayesian inference for linear and logistic regression." *arXiv preprint arXiv:1310.5438* (2013).